L Number	Hits	Search Text	DB	Time stamp
15	11	6192401.URPN.	USPAT	2004/08/19 15:38
16	431	shut adj down with (node)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 15:40
17	35	shut adj down with (node) same cluster	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 15:40
20	287	shut adj down with (node)	USPAT	2004/08/19 15:41
21	9	shut adj down with (node) with (weight$3 or value or priority or significance)	USPAT	2004/08/19 15:42
22	2	("6151688").PN.	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 15:42
-	482	dynamic$8 with priority with (application or program or software)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 10:02
-	106	dynamic$8 near3 priority near3 (application or program or software)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/18 18:23
-	56	(dynamic$8 near3 priority near3 (application or program or software)) and ((@ad < "19990528") or (@prad < "19990528") or (@rlad < "19990528"))	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/18 18:23
-	2	("6279032").PN.	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 14:58
-	6	6279032.URPN.	USPAT	2004/08/19 10:02
-	85	(cluster or partition) with node with (vote or status or state or priority or weight) with (application or program or process)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 10:52
-	44	((cluster or partition) with node with (vote or status or state or priority or weight) with (application or program or process)) and ((@ad < "19990528") or (@prad < "19990528") or (@rlad < "19990528"))	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 14:58
-	12	(cluster or partition) with node with (vote or priority or weight) with (application or program)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 10:53
-	14	(cluster or partition) same node with (vote or priority or weight) with (application or program)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 11:03

-	2	(cluster or partition) same (vote or priority or weight) with (application or program) same load adj balanc$3	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 11:06
-	9	("5852717" | "6098093" | "6119153" | "6178461" | "6185608" | "6253234" | "6332163" | "6377991" | "6480865").PN.	USPAT	2004/08/19 11:05
-	0	6697849.URPN.	USPAT	2004/08/19 11:06
-	2	(cluster or partition) same (vote or priority or weight) with (application or program) same arbitrat$4	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 11:08
-	6	(cluster or partition) and (node same (vote or priority or weight) with (application or program) same arbitrat$4)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 11:10
-	34	(cluster or partition) and ((vote or priority or weight) with (application or program) same arbitrat$4)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 11:14
-	258	(cluster or partition) and ((vote or priority or weight) with (application or program) same member$4)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 11:14
-	53	(cluster or partition) and ((vote or priority or weight) with (application or program) same membership)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 11:15
-	5	(cluster or partition) and ((vote or priority or weight) with (application or program) same membership)	USPAT	2004/08/19 11:20
-	9	(cluster or partition) and ((vote or priority or weight) with (application or program) same member$4 with (determin$6 or resolv$6))	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 11:17
-	19	(cluster or partition) same ((vote or priority or weight) with (application or program) with (calculat$5 or determin$6 or identify$4))	USPAT	2004/08/19 11:21
-	0	partition adj in adj space	USPAT	2004/08/19 11:22
-	1	partition-in-space	USPAT	2004/08/19 11:22
-	16	("4564903" | "4843541" | "5220654" | "5257368" | "5257379" | "5301323" | "5317739" | "5473773" | "5526484" | "5530860" | "5561809" | "5564040" | "5675739" | "5706432" | "5881284" | "5896520").PN.	USPAT	2004/08/19 11:25
-	11	6192401.URPN.	USPAT	2004/08/19 11:28
-	11	6192401.URPN.	USPAT	2004/08/19 11:34
-	37	("5280627" | "5404527" | "5553239" | "5555375" | "5659748" | "5673384" | "5727206" | "5754821" | "5781910" | "5828876" | "5828889" | "5828961" | "5835784" | "5892913" | "5893086" | "5909540" | "5917998" | "5918229" | "5927050" | "5940838" | "5946686" | "5948109" | "5974547" | "5996075" | "5999712" | "5999978" | "6014669" | "6108699" | "6108781" | "6192401" | "6301462" | "6311217" | "6314526" | "6360331" | "6363495" | "6427163" | "6438705" | "2001/0014097").PN.	USPAT	2004/08/19 11:35
-	67	Oracle adj Parallel adj Server	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 11:49

-	38	Oracle adj Parallel adj Server same (cluster or partition)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 11:49
-	17	Oracle adj Parallel adj Server same (cluster or partition)	USPAT	2004/08/19 11:49
-	11	6272544.URPN.	USPAT	2004/08/19 13:54
-	11	6272544.URPN.	USPAT	2004/08/19 13:55
-	6	6279032.URPN.	USPAT	2004/08/19 14:46
-	193	(cluster or partition) same (node) with weight$2	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 14:58
-	101	(cluster or partition) same (node) with weight$2	USPAT	2004/08/19 15:16
-	84	((cluster or partition) same (node) with weight$2) and ((@ad < "19990528") or (@prad < "19990528") or (@rlad < "19990528"))	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 15:06
-	914	(dynamic$4 or run adj time) with (creat$4 or generat$4) with (class)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 15:13
-	192	(dynamic$4 or run adj time) with (creat$4 or generat$4) adj2 (class)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 15:13
-	320	(dynamic$4 or run adj time) with (creat$4 or generat$4) near2 (class)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 15:13
-	278	(dynamic$4 or run adj time) near5 (creat$4 or generat$4) near2 (class)	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 15:13
-	3	(dynamic$4 or run adj time) near5 (creat$4 or generat$4) near2 (class) with generic$4	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 15:14
-	2	("5481718").PN.	USPAT; US-PGPUB; EPO; DERWENT; IBM_TDB	2004/08/19 15:14
-	7	(cluster or partition) same (node) with weight$2 with (application or program or software)	USPAT	2004/08/19 15:17

CiteSeer Find: node weight and cluster and failov Documents Citations

Searching for **node weight and cluster and failover**.
Restrict to: Header Title Order by: Expected citations Hubs Usage Date Try: Google (CiteSeer) Google (Web) CSB DBLP
No documents match Boolean query. Trying non-Boolean relevance query.
500 documents found. **Order: relevance to query.**

Efficient Support for P-HTTP in Cluster-Based Web Servers - Aron, Druschel, Zwaenepoel (1999) (Correct) (10 citations)
of HTTP/1.1 requests among the back-end **nodes** of a **cluster** server. A trace-driven simulation
www.cs.rice.edu/~aron/papers/phttp-lard.ps

Using Constraints To Achieve Stability In Automatic Graph.. - Böhringer, Paulisch (1990) (Correct)
1 INTRODUCTION Graphs, consisting of a set of **nodes** and a set of edges, are one of the most
B" or "**node** C is the top neighbor of **node** D"**Clusters**: Gather a group of **nodes** together to a "**cluster**"
www.ee.washington.edu/faculty/karl/Publications/PS/SIGCHI90.ps.gz

Correction of a Memory Management Method for Lock-Free Data.. - Michael, Scott (1995) (Correct) (7 citations)
care. Many lock-free algorithms require deleted **nodes** not to be reused until no active pointers point
hypatia.dcs.qmw.ac.uk/data/edu/cs.rochester.edu/systems/95.tr599.Memory_management_for_lock-free_data_structures.ps.gz

The Application Of The Supercondensed Tlm Node To.. - Trenkic.. (Correct)
Limited The Application Of The SupercondensedTlm **Node** To Propagation Problems In Inhomogeneous
nmle.eee.nott.ac.uk/~vmt/preprints/compel95.ps.gz

Surface Brightness Evolution of Cluster and Field Galaxies - Schade (Correct)
Surface Brightness Evolution of **Cluster** and Field Galaxies DavidSchade 1 National
www.dao.nrc.ca/science/preprint/97/9702.ps.gz

A Unified Network-based Approach for Online Recognition of.. - Lee, Kim (Correct)
a word network, a path spanning from the initial **node** to the final **node**, possibly with a number of
ai.kaist.ac.kr/~joony/ps/IWFHR_96.ps

Performance Evaluation and Modeling of MPI Communications .. - Folino, Spezzano, Talia (Correct)
computer. It consists of Sparc based processing **nodes** running the Solaris operating system on each
operating system on each **node**, so it resembles a **cluster** of workstations connected by a fast network.
isi-cnr.deis.unical.it:1080/~talia/hpcn98.ps

The Beehive Cluster System - Aman Singla (Correct)
Box 2 CPUs 1 CPU threads in a single process parent **node** Interconnection Network Figure 1: The Beehive
multiple units of CPU scheduling called lwp (light-**weight** process) within the same address space. It also
The Beehive **Cluster** System Aman Singla Umakishore Ramachandran
www.cc.gatech.edu/computing/Architecture/Beehive/../papers/beehive.ps.gz

Identification of a high redshift cluster in the field of.. - Pelló, Miralles (1996) (Correct)
redshift intervals have been cumulated with a **weight** of 1. A significant excess of galaxies at a mean
20.3.1996 Identification of a high redshift **cluster** in the field of Q2345007 through deep BRIJK'
www.astr.tohoku.ac.jp/~miralles/AA_2345_3nov95.ps.gz

Using PVM 3.0 to Run Grand Challenge Applications on.. - Dongarra, Geist.. (1992) (Correct)
This is the ability to run PVM applications on the **nodes** of several different distributed memory
is to exploit the aggregate power and memory of **clusters** of computers sitting on local area networks.
ftp.netlib.org/ncwn/siam93-pvmgc.ps

Topic Detection and Tracking Pilot Study - Allan, Carbonell, Doddington.. (1998) (Correct) (3 citations)
a speech recognizer with 100 underlying "single **node**" models (corresponding to the topics)each of
result from adding the feature and adjusting its **weight** to the best value. After calculating the gain of
is to group the stories in the study corpus into **clusters**, where each **cluster** represents an event and
www.cs.cmu.edu/~yiming/papers.yy/tdt1-final-report.ps

Uniform Reconstruction of Gaussian Processes - Müller-Gronbach, Ritter (1995) (Correct) (1 citation)
each realization of X the error is defined in a **weight**ed supremum norm the overall error of a
ftp.math.fu-berlin.de/pub/math/publ/pre/1995/pr-a-95-26.ps.Z

Optimized Software Synthesis for Digital Signal.. - Jürgen Teich.. (1998) (Correct) (3 citations)
consists of a directed graph in which the **nodes** represent computations and the arcs specify the
component into an acyclic graph that consists of **clusters** of smaller strongly connected components,
ftp.tik.ee.ethz.ch/pub/people/zitzler/TZB1998a.ps.gz

SemQuery: Semantic Clustering and Querying on.. - Sheikholeslami.. (1998) (Correct) (2 citations)
12) where the summation is done over all input **nodes** j, and t k is the desired value of output o k
image content-based retrieval systems use a **weight**ed Euclidean distance to combine the similarity
SemQuery: Semantic **Cluster**ing and Querying on Heterogeneous Features for
www.rit.edu/~wcceec/./papers/tkde-semantic.ps

Daily management of an earth observation satellite.. - Lemaitre, Verfaillie (Correct)
for the next day, each one associated with a **weight** reflecting its importance, 2) a set of
ftp.cert.fr/pub/lemaitre/Papers/97-ILOG.ps

Pitch Determination Considering Laryngealization.. - Niemann, Denzler, .. (1994) (Correct) (2 citations)
I =0 having i =0 1 M 0 \Gamma 1 input **nodes** n 0) i for signal values x i ,ffl the hidden
for each layer and each **node** per layer compute the **weight**ed sum of input values y I1) j =M I \Gamma1
www5.informatik.uni-erlangen.de/TeX/Literatur/ps-dir/1994/Niemann94:PDCa.ps.gz

F3.733e+05>J. Phys.: Condens. Matter 8 .. - Orlando Azavant (Correct)
Matter 8 (1996) 1123-1133. Printed in the UK **Cluster** and supercell calculations for carbon-doped
www.tcm.phy.cam.ac.uk/~mdt26/downloads/cinsi.ps.gz

A Partial Approach to the Problem of Deadlocks in.. - Tricas.. (1998) (Correct)
generic way, elements belonging to P [T are called **nodes**. The set W :P \Theta T)T \Theta P)IN
then we say that there is an arc from x to y, with **weight** or multiplicity W (x y)Ordinary nets are those
www.cps.unizar.es/~ftricas/GISIRR9705.ps.gz

Solving Small TSPs with Constraints - Caseau, Laburthe (1997) (Correct) (20 citations)
a technique of choice for solving small (up to 30 **nodes**) traveling salesman problems. These techniques
some edges early on. Indeed, when an edge ij has a **weight** much higher than the best assignment for i or j,
geographic map and when the cities are grouped in **clusters** on this map, neither the bound SN nor the
www.dmi.ens.fr/users/laburthe/papers/iclp97.ps.gz

Optimal Emulation of Meshes on Meshes of Trees - Achilles (1995) (Correct) (5 citations)
mesh. In this paper, we present an emulation of N-**node** meshes on O(N)**node** meshes of trees with constant
liinwww.ira.uka.de/~achilles/MOT.EUROPAR.95.ps.gz

First 20 documents Next 20

CiteSeer Find: node weight and cluster Documents Citations

Searching for **node weight and cluster**.
Restrict to: Header Title Order by: Expected citations Hubs Usage Date Try: Google (CiteSeer) Google (Web) CSB DBLP
6 documents found. **Order: number of citations.**

WCA: A Weighted Clustering Algorithm for Mobile Ad hoc.. - Chatterjee, Das, Turgut (2001) (Correct) (11 citations)
heuristic, ii) Lowest-ID heuristic, and (iii) **Node-Weight** heuristic. The Lowest-ID and the
Preprint 0 (2001) 1 WCA: A Weighted **Clustering** Algorithm for Mobile Ad hoc Networks
In this paper, we propose an on-demand distributed **clustering** algorithm for multi-hop packet radio
crewman.uta.edu/~mainak/papers/cluster.ps

A Weight Based Distributed Clustering Algorithm for.. - Chatterjee, Das, Turgut (1970) (Correct) (3 citations)
heuristic (ii) Lowest-ID heuristic and (iii) **Node-Weight** heuristic. In the assumed graph model of the
A Weight Based Distributed **Clustering** Algorithm for Mobile Ad hoc Networks
Abstract. In this paper, we propose a distributed **clustering** algorithm for a multi-hop packet radio
crewman.uta.edu/~mainak/papers/hipc00.ps

An On-Demand Weighted Clustering Algorithm (WCA) for Ad.. - Chatterjee, Das, Turgut (2000) (Correct) (1 citation)
heuristic (ii) Lowest-ID heuristic and (iii) **Node-Weight** heuristic. In the assumed graph model of the
An On-Demand Weighted **Clustering** Algorithm (WCA) for Ad hoc Networks Mainak
Abstract-We consider a multi-**cluster**, multi-hop packet radio network architecture for
crewman.uta.edu/~mainak/papers/globecom00.ps

A New Heuristic for Vehicle Routing with Narrow Time Windows - Hamacher, Moll (1996) (Correct) (1 citation)
clustering procedure where a tree with multiple **node weights** is divided into subtrees. Upper bounds
planned stops. The algorithm is based on a **clustering** procedure where a tree with multiple node
The algorithm is divided into two parts. In the **clustering** step the customers are partitioned into
ftp.mi.uni-koeln.de/pub/paper/postscript/zpr96-239.ps

Cluster Computing 5, 193--204, 2002 - Wca Weighted Clustering (2002) (Correct)
heuristic, ii) Lowest-ID heuristic, and (iii) **Node-Weight** heuristic. The Lowest-ID and the
Cluster Computing 5, 193-204, 2002 2002 Kluwer
Manufactured in The Netherlands. WCA: A Weighted **Clustering** Algorithm for Mobile Ad Hoc Networks
crewman.uta.edu/~turgut/published_papers/WCA-2002.pdf

Adaptive Load Distribution in Strong Dynamic Heterogeneous Systems - Rips (Correct)
same way using the average computation speed, a **node weight** is determined. Edge weight and **node weight** can
speed, a **node weight** is determined. Edge weight and **node weight** can be defined as: Definition 1 Let
(FCP) which runs on heterogeneous workstation **clusters** using PVM (Parallel Virtual Machine [7]The
www.hlrs.de/people/rips/Paper/mpcs96.ps

Try your query at: Google (CiteSeer) Google (Web) CSB DBLP

Terms used **node cluster weight failover** Found **11** of **141,345**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 11 of 11

Relevance scale

1 Fast cluster failover using virtual memory-mapped communication
Yuanyuan Zhou, Peter M. Chen, Kai Li
May 1999 **Proceedings of the 13th international conference on Supercomputing**
Full text available: pdf(1.45 MB) Additional Information: full citation, references, citings, index terms

2 Analysis and implementation of software rejuvenation in cluster systems
Kalyanaraman Vaidyanathan, Richard E. Harper, Steven W. Hunter, Kishor S. Trivedi
June 2001 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2001 ACM SIGMETRICS international conference on Measurement and modeling of computer systems**, Volume 29 Issue 1
Full text available: pdf(983.05 KB) Additional Information: full citation, abstract, references, citings

Several recent studies have reported the phenomenon of "software aging", one in which the state of a software system degrades with time. This may eventually lead to performance degradation of the software or crash/hang failure or both. "Software rejuvenation" is a pro-active technique aimed to prevent unexpected or unplanned outages due to aging. The basic idea is to stop the running software, clean its internal state and restart it. In this paper, we discuss software rejuvenation as applied to ...

3 Cluster-based scalable network services
Armando Fox, Steven D. Gribble, Yatin Chawathe, Eric A. Brewer, Paul Gauthier
October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles**, Volume 31 Issue 5
Full text available: pdf(2.42 MB) Additional Information: full citation, references, citings, index terms

4 Bayeux: an architecture for scalable and fault-tolerant wide-area data dissemination
Shelley Q. Zhuang, Ben Y. Zhao, Anthony D. Joseph, Randy H. Katz, John D. Kubiatowicz
January 2001 **Proceedings of the 11th international workshop on Network and operating systems support for digital audio and video**
Full text available: pdf(272.26 KB) Additional Information: full citation, abstract, references, citings, index terms

The demand for streaming multimedia applications is growing at an incr edible rate. In this paper, we propose Bayeux, an efficient application-level multicast system that scales to arbitrarily large receiver groups while tolerating failures in routers and network links. Bayeux also includes specific mechanisms for load-balancing across replicate root nodes and more efficient bandwidth consumption. Our simulation results indicate that Bayeux

maintains these properties while keeping transmi ...

5 BuddyCache: high-performance object storage for collaborative strong-consistency applications in a WAN

Magnus E. Bjornsson, Liuba Shrira

November 2002 **ACM SIGPLAN Notices , Proceedings of the 17th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications**, Volume 37 Issue 11

Full text available: pdf(269.18 KB) Additional Information: full citation, abstract, references, index terms

Collaborative applications provide a shared work environment for groups of networked clients collaborating on a common task. They require strong consistency for shared persistent data and efficient access to fine-grained objects. These properties are difficult to provide in wide area networks because of high network latency.*BuddyCache* is a new transactional caching approach that improves the latency of access to shared persistent objects for collaborative strong-consistency applications i ...

Keywords: cooperative caching, fault-tolerance, fine-grain sharing, object storage systems, transactions, wide-area network

6 Industrial sessions: middle-tier caching: Middle-tier database caching for e-business

Qiong Luo, Sailesh Krishnamurthy, C. Mohan, Hamid Pirahesh, Honguk Woo, Bruce G. Lindsay, Jeffrey F. Naughton

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available: pdf(1.20 MB) Additional Information: full citation, abstract, references, citings, index terms

While scaling up to the enormous and growing Internet population with unpredictable usage patterns, E-commerce applications face severe challenges in cost and manageability, especially for database servers that are deployed as those applications' backends in a multi-tier configuration. Middle-tier database caching is one solution to this problem. In this paper, we present a simple extension to the existing federated features in DB2 UDB, which enables a regular DB2 instance to become a DBCache wi ...

7 Managing energy and server resources in hosting centers

Jeffrey S. Chase, Darrell C. Anderson, Prachi N. Thakar, Amin M. Vahdat, Ronald P. Doyle

October 2001 **ACM SIGOPS Operating Systems Review , Proceedings of the eighteenth ACM symposium on Operating systems principles**, Volume 35 Issue 5

Full text available: pdf(1.61 MB) Additional Information: full citation, abstract, references, citings, index terms

Internet hosting centers serve multiple service sites from a common hardware base. This paper presents the design and implementation of an architecture for resource management in a hosting center operating system, with an emphasis on *energy* as a driving resource management issue for large server clusters. The goals are to provision server resources for co-hosted services in a way that automatically adapts to offered load, improve the energy efficiency of server clusters by dynamically res ...

8 Parallel and distributed systems and networking: A parallel index for semistructured data

Brian F. Cooper, Neal Sample, Moshe Shadmon

March 2002 **Proceedings of the 2002 ACM symposium on Applied computing**

Full text available: pdf(638.30 KB) Additional Information: full citation, abstract, references, index terms

Database systems are increasingly being used to manage semistructured data, which may not have a fixed structure or set of relationships between data items. Indexes which use tree structures to manage semistructured data become unbalanced and difficult to parallelize due to the complex nature of the data. We propose a mechanism by which an

unbalanced *vertical* tree is managed in a balanced way by additional layers of *horizontal* index. Then, the vertical tree can be partitioned among ...

9 Session III: Mobility Management in multimedia networks: Mobility support in unified communication networks

Helen J. Wang, Randy H. Katz

July 2001 **Proceedings of the 4th ACM international workshop on Wireless mobile multimedia**

Full text available: pdf(1.11 MB) Additional Information: full citation, abstract, references, index terms

Rapid advances in communication networks and device technologies have enabled people with powerful means of communications. It is common for any individual to be associated with a number of heterogeneous communication devices (such as phones, pagers, PDAs) or a variety of applications (such as e-mail, instant messaging, or chat-rooms). This phenomenon has spurred a great demand for *unifie dcommunication* [20] services which integrate one's various communication mechanisms in a meaningful a ...

10 Fair queuing for aggregated multiple links

Josep M. Blanquer, Banu Özden

August 2001 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications**, Volume 31 Issue 4

Full text available: pdf(184.49 KB) Additional Information: full citation, references, citings, index terms

11 Tomography: Tomography-based overlay network monitoring

Yan Chen, David Bindel, Randy H. Katz

October 2003 **Proceedings of the 2003 ACM SIGCOMM conference on Internet measurement**

Full text available: pdf(259.22 KB) Additional Information: full citation, abstract, references, index terms

Overlay network monitoring enables distributed Internet applications to detect and recover from path outages and periods of degraded performance within seconds. For an overlay network with n end hosts, existing systems either require $O(n^2)$ measurements, and thus lack scalability, or can only estimate the latency but not congestion or failures. Unlike other network tomography systems, we characterize end-to-end losses (this extends to any additive metrics, including latency) rather tha ...

Keywords: network measurement and monitoring, network tomography, numerical linear algebra, overlay networks

Results 1 - 11 of 11